Coverage Path Planning for Robotic De-mining

Howie Choset
Department of Mechanical Engineering
Carnegie Mellon
Pittsburgh, PA 15213
Phone: (412) 268-2495 Fax: (412) 268-3348 E-mail: choset@cs.cmu.edu

Grant Number: N00014-99-1-0478 (YIP)
http://voronoi.sbp.ri.cmu.edu/

LONG-TERM GOALS

This work develops key technologies for de-mining a target region, whether it is in shallow water, on land, or at the interface between them. The crucial first step of de-mining is finding the mines. Searching for mines is a dangerous and expensive task. The use of robots immediately bypasses the danger, reduces the cost, and potentially speeds the process. In de-mining, a robot must pass a mine-detecting sensor over all points in the region that might conceal a mine. To do this, the robot must traverse a carefully planned path through the target region.

OBJECTIVES

Conventional path planners are inadequate for de-mining because they only produce paths between two points and pay no attention to the intervening area. Coverage-path planning, as its name suggests, specifically emphasizes the space swept out by the robot's sensor. Integrating the robot's footprint (detector range) along the coverage path yields an area identical to that of the target region.

Most coverage-path planners are rudimentary, at best, because they rely on heuristics. As merely rules of thumb, heuristics cannot provide any guarantees. The proposed work develops complete techniques, provably ensuring that the detector passes over all points in the region. The guarantee of completeness removes a potential source of uncertainty that the robot has done its job thoroughly, which is essential in applications such as minesweeping.

Successfully implementing any path planner in a large-scale region requires that a robot can determine its actual position. The algorithms under development for this work also address this problem by automatically encoding environmental characteristics that allow the robot to correct internal positioning errors. This localization capability can be especially useful in applications where current practice uses explicitly added landmarks, such as transponders.

Transponder compromise the covert nature of the de-mining mission and thus eliminating them would be of great utility. In searching for unexploded ordnance (UXO), transponders, in the form of GPS, can be used but are quite costly. In conjunction with the Naval Explosive Ordnance Disposal (EOD) Technology Division, we are developing low cost technologies to provide position information to robots covering target (and potentially unknown) regions.

In many scenarios, time and other resources may not be in abundance, and thus if the planner has access to a probabilistic map of mine locations, it can guide opportunistically the robot. The robot can use this approach for rapid lane clearing.

APPROACH

To ensure complete coverage, this work employs a cellular decomposition, which divides the target region into sub-regions, called cells, such that simple back-and-forth motions, like that of an ox plowing a field, suffices to cover the cell. The planner achieves complete coverage simply by ensuring that the robot identifies and visits each cell in the region. Prior work includes developing a decomposition called the boustrophedon decomposition and this effort has developed a procedure to implement the boustrophedon decomposition in an unknown environment, i.e., allow for coverage of unknown regions.

This work uses two approaches to address the positioning problem while covering large-scale regions: algorithmic (software) and technological. The boustrophedon decomposition solves this problem by automatically encoding environmental characteristics that allow the robot to correct internal positioning errors. Each cell in the decomposition has a characteristic or sensor signature that may or may not be unique to the cell. Adjacent cells share an edge between them. By understanding the adjacency relationship among cells, the robot can locate itself; this aspect of the work will be addressed in the second and third years of the proposed effort.

Currently, we are working with the Naval Explosive Ordnance Disposal Division developing an inexpensive outdoor mobile platform with inexpensive positioning technology to bypass GPS units that are expensive and cumbersome to integrate. We are developing two systems: one using a linear encoder and the other a camera that looks at well-marked posts. A linear encoder is nothing more than fishing wire on a high-precision wheel. The wire is fastened to a post. We use two of these wheel-wire systems with posts that are a fixed and known distance a part from each other. As the robot moves about in its environment, the lengths of these wires provide an accurate location of the robot.

Granted that these wires on the robot will be a bit cumbersome to deploy, but this system will constitute the first out-door mobile robot that accurately covers an unknowns space, on the order of 50 by 50 feet. This system is far less expensive and more accurate than GPS. Unfortunately, these wires may be too and thus we will put inexpensive cameras that look at hand-engineered posts to replace these wires. With accurate positioning and navigation capabilities, we can ensure complete coverage of a target region on a real robot in a real UXO search scenario.

Finally, we have made significant inroads towards developing probabilistic techniques to guide the robot to opportunistically exploit any form of a priori information, if it exists. In such scenarios, the robot can cover a small percentage of the target region while locating a large percentage of the UXO's and mines. For now, we are assuming a minefield that has been laid out according to a pattern that was suggested to us by Chris Debolt at the Naval Explosive Ordnance Disposal Division. This minefield can be described by six parameters: inter-row spacing, inter-column spacing, an offset, the origin of the minefield, and its orientation. We use a Bayesian-based approach to estimate the parameters and direct the robot's motion.

WORK COMPLETED

This year we finished all of the derivations required for complete sensor based coverage of an unknown region and implemented it on an indoor mobile robot equipped only with simple and inexpensive range sensors. We submitted a paper to Tuan Nguyen at the Naval Explosive Ordnance Disposal Technology Division for his review and after receiving feedback from him, we will submit this paper to the International Journal of Robotics Research. This algorithm differs from prior coverage work in that, barring positioning error, it is guaranteed to cover a target region and does not make any unrealistic assumptions like the target region can be represented by a fine resolution grid (of pixels). Experiments with the mobile robot clearly point to the need for good positioning algorithms.

Our collaborators at the Naval Explosive Ordnance Disposal Technology Division want us to demonstrate our algorithms outdoors on an inexpensive mobile robot that has a low computational, memory, and power budget. Ideally, we would use the robots that they are developing in house, but they were not fully developed, so we began our own inexpensive outdoor mobile robot effort at Carnegie Mellon. In doing this, we have developed an insight on how out-door breadbox sized robots should be designed and built. In fact, it is our hope that the Tuan Nguyen's internal development will benefit from ours, and visa versa. At the very least, we will be developing a range-sensing module that will be compatible with the Naval Ordnance Technology Division's robot.

Already, we have developed the first iteration of this robot covering a grassy, bumpy, and wet 30 foot by 30 foot area outdoors. We designed and built the linear encodes, described above, that enabled us to accurately guide the robot outdoors. Unfortunately, we were limited to 30 feet because of limitations of the linear encoders. We spent some time trying to improve this limit, but abandoned the approach for a vision-based system. The vision-based system will use engineered landmarks that are placed in the environment. Already, we have developed and integrated the hardware for this system, as well as some basic software development. We expect to have this system working by December 2000 for Tuan Nguyen's second visit to Carnegie Mellon.

Finally, we have developed a Bayesian-based algorithm that estimates the parameters of a minefield whose mines are laid out according the pattern described above. We have also begun writing software that implements these algorithms. We intend to finish the software within the next month, run simulations on it, and then port the software to the outdoor de-mining robot this year.

We have had three demonstrations of our work this year. The first was to Tom Swean and Chris O'Donnell in Panama City (video only), the second was to Tom Swean, Chris Debolt, and Tuan Nguyen in Washington, DC (video and presentation), and the third was to Tuan Nguyen at Carnegie Mellon University (presentations and live demonstrations). Everything summarized in this report was demonstrated at the Carnegie Mellon demonstration. Tuan Nguyen will return to Carnegie Mellon in December.

IMPACT/APPLICATIONS

This technology will impact all de-mining and UXO searching operations. Not only will it remove people from danger, it will expedite these missions with more accuracy.

TRANSITIONS

Currently, the Principal Investigator's group is working with the Naval EOD Technology Division in developing an outdoor UXO searching mobile base for the purposes of transporting algorithms from the proposed effort onto their base. The Carnegie Mellon University (CMU) base will be as compatible as possible with the Naval EOD Technology Division base to allow for easy software transition.

RELATED PROJECTS

Already, we are working with the Naval EOD Technology Division at Indian Head who is running the Basic UXO Gathering System (BUGS) program that uses many robots to search for UXO's. The work in this effort will directly carry over to the BUGS project. Foster-Miller develops technologies for Autonomous Underwater Vehicles (AUV's) using synthetic aperture sonar to look for features, such as mines, under water. The positioning algorithms that we will work on in year two will provide an efficient algorithm to allow the robot to use the features from the sonar map to position itself relative to a prior known location.

We have also started to apply our coverage work to car painting with Ford. Today, a car painting specialist requires three-to-five months to program a car-painting robot to cover the surface of a car for the painting application. The car-painting specialist does not use any standard path planning technique to "semi-automatically program" the robot to paint the car, but instead the specialist "manually programs" the robot in an attempt to ensure full coverage of the car body. This process requires the programmer to position (possibly in an off-line simulation programming environment) the robot at a large number (hundreds) of points along its path over the car body

We are lifting our planar coverage work into three dimensions. By itself, this is challenging, but the added difficulty is that the robot must deposit a uniform thickness of paint on the surface of the car. With de-mining, the robot can pass over the same points multiple times without affecting performance; with car painting, the robot must be incredibly careful when passing over already-painted points so as to ensure uniform paint thickness. We believe that this work will have a Naval application of inspecting the hulls of ships for cracks and other sources of failure.